

RISK, AMBIGUITY, AND THE SAVAGE AXIOMS

Daniel Ellsberg

August 1961

A familiar, ongoing pattern of activity may be subject to considerable uncertainty, but this uncertainty is more apt to appear in the form of "risk"; the relation between given states of nature is known precisely, and although the random variation in the state of nature which "obtains" may be considerable, its stochastic properties are often known confidently and in detail. (Actually, this confidence may be self-deceptive, based on ignoring some treacherous possibilities; nevertheless, it commonly exists). In contrast, the ambiguities surrounding the outcome of a proposed innovation, a departure from current strategy, may be much more noticeable. Different sorts of events are relevant to its outcome, and their likelihoods must now be estimated, often with little evidence or prior expertise; and the effect of a given state of nature upon the outcome of the new action may itself be in question. Its variance may not appear any higher than that of the familiar action when computed on the basis of "best estimates" of the probabilities involved, yet the meaningfulness of this calculation may be subject to doubt. The decision rule ^{not recommended; or, explaining observed, relevant} discussed will not preclude choosing such an act, but it will definitely bias the choice away from such ambiguous ventures and toward the strategy with "known risks." Thus the rule is "conservative" in a sense more familiar to everyday conversation than to statistical decision theory; it may often favor traditional or current strategies, even perhaps at high risk, over innovations whose consequences are undeniably ambiguous. This property may recommend it to some, discredit it with others (some of whom might prefer to reverse the rule, to emphasize the more hopeful possibilities in ambiguous situations); it does not seem irrelevant to one's attitude toward the behavior.

In the equivalent formulation in terms of y_x^{\min} and y^0 , the subject

Aug '61

above could be described "as though" he were assigning weights to the respective payoffs of actions II and III, whose expected values are ambiguous, as follows (assuming $y^0 = (1/3, 1/3, 1/3)$ in each case):

	y_x^{\min}	$\rho \cdot y^0 + (1-\rho) y_x^{\min}$
II	$(\frac{1}{3}, 0, \frac{2}{3})$	$(\frac{1}{3}, \frac{1}{12}, \frac{7}{12})$
III	$(\frac{1}{3}, \frac{2}{3}, 0)$	$(\frac{1}{3}, \frac{7}{12}, \frac{1}{12})$

Although the final set of weights for each set of payoffs resemble probabilities (they are positive, sum to unity, and represent a linear combination of two probability distributions), they differ for each action, since y_x^{\min} will depend on the payoffs for x and will vary for different actions. If these weights were interpreted as "probabilities," we would have to regard the subject's subjective probabilities as being dependent upon his payoffs, his evaluation of the outcomes. Thus, this model would be appropriate to represent cases of true pessimism, or optimism/wishfulness (with y_x^{\max} substituting for y_x^{\min}). However, in this case we are assuming conservatism, not pessimism; our subject does not actually expect the worst, but he chooses to act "as though" the worst were somewhat more likely than his best estimates of likelihood would indicate. In either case, he violates the Savage axioms; it is impossible to infer from the resulting behavior a set of probabilities for events independent of his payoffs. In effect, he "distorts" his best estimates of likelihood, in the direction of increased emphasis on the less favorable outcomes and to a degree depending on ρ , his confidence in his best estimate.²⁰

Not only does this decision model account for "deviant" behavior in a particular, ambiguous situation, but it covers the observed shift in a